

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

Bluetooth—The universal radio interface for *ad hoc*, wireless connectivity

Jaap Haartsen

V8039

XP-000783249

p. 110-117 = (8)

p.d. 1998

Bluetooth is a universal radio interface in the 2.45 GHz frequency band that enables portable electronic devices to connect and communicate wirelessly via short-range, *ad hoc* networks. Each unit can simultaneously communicate with up to seven other units per piconet. Moreover, each unit can simultaneously belong to several piconets.

Bluetooth technology—which apart from Ericsson, has gained the support of Nokia, IBM, Toshiba, Intel and many other manufacturers—eliminates the need for wires, cables and connectors for and between cordless or mobile phones, modems, headsets, PDAs, computers, printers, projectors, local area networks, and so on, and paves the way for new and completely different devices and applications.

Before guiding us through frequency-hopping technology and the channel, packet and physical-link definitions that characterize the Bluetooth air interface, the author briefly describes the conditions that led up to the development of Bluetooth. He then acquaints us with the networking aspects of Bluetooth technology, describing piconets and scatternets, connection procedures, and inter-piconet communication.

Imagine a cheap, power-efficient radio chip that is small enough to fit inside any electronic device or machine, that provides local connectivity, and that creates a (worldwide) micro-scale web. What applications might you use it in?

In 1994, Ericsson Mobile Communications AB in Lund, Sweden, initiated a study to investigate the feasibility of a low-power, low-cost radio interface between mobile phones and their accessories. The intention

was to eliminate cables between phones and PC cards, wireless headsets, and so forth. The study was part of a larger project that investigated multi-communicators connected to the cellular network via cellular telephones. The last link in the connection between a communicator and the cellular network was a short-range radio link to the phone—thus, the link was called the multi-communicator link or MC link. As the MC link project progressed, it became clear that there was no limit to the kinds of application that could use a short-range radio link. Cheap, short-range radios would make wireless communication between portable devices economically feasible.

Current portable devices use infrared links (IrDA) to communicate with each other. Although infrared transceivers are inexpensive, they

- have limited range (typically one to two meters);
- are sensitive to direction and require direct line-of-sight;
- can in principle only be used between two devices.

By contrast, radios have much greater range, can propagate around objects and through various materials, and connect to many devices simultaneously. What is more, radio interfaces do not require user interaction.

In the beginning of 1997, when designers had already begun work on an MC link

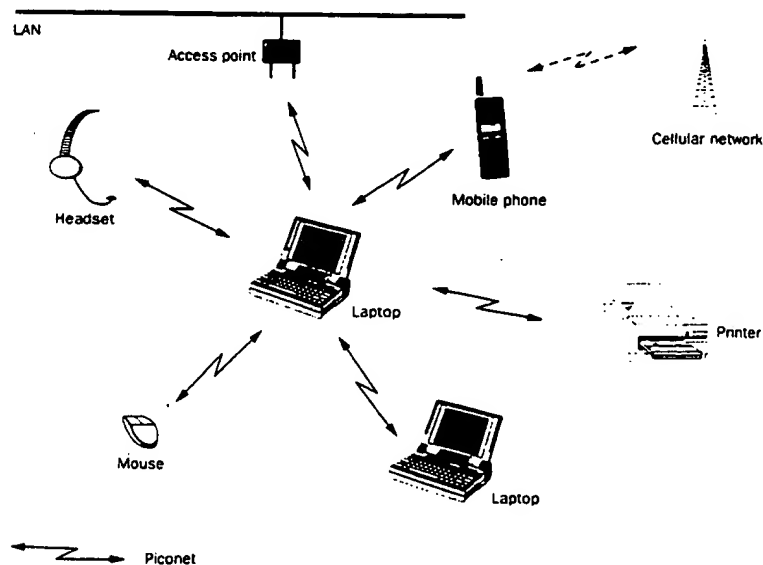


Figure 1
User model with local wireless connectivity.
Applications envisioned for the near future.

transceiver chip. Ericsson approached other manufacturers of portable devices to raise interest in the technology—for the system to succeed, a critical mass of portable devices must use the short-range radio. In February 1998, five promoters—Ericsson, Nokia, IBM, Toshiba and Intel—formed a special interest group (SIG). The idea was to achieve a proper mix of business areas: two market leaders in mobile telephony, two market leaders in laptop computing, and a market leader in core, digital-signal-processor (DSP) technology. On May 20 and 21, 1998, the Bluetooth consortium announced itself to the general public from London, England; San Jose, California; and Tokyo, Japan. Since then, several companies have joined as adopters of the technology (Box B).

The purpose of the consortium is to establish a *de facto* standard for the air interface and the software that controls it, thereby ensuring interoperability between devices of different manufacturers. The first products to use MC link technology will emerge at the end of 1999 in mobile phones, notebook computers and accessories (Figure 1).

Box A Abbreviations

ACL	Asynchronous connectionless
ARQ	Automatic retransmission query
CVSD	Continuous variable slope delta
DSP	Digital signal processor
FEC	Forward error correction
FH	Frequency hop
FSK	Frequency shift keying
HEC	Header error correction
HPC	Handheld personal computer
IrDA	Infrared Data Association
ISM	Industrial Scientific Medical
MAC	Media access control
MC	Multicomunicator
PC	Personal computer
PDA	Personal digital assistant
RF	Radio frequency
SCO	Synchronous connection-oriented
SIG	Special interest group
TDD	Time division duplex
TDM	Time division multiplex

Box B

The Bluetooth consortium—promoters and adopters

The promoters of the Bluetooth* consortium formed a special interest group (SIG) at Ericsson Inc., Research Triangle Park, North Carolina, on February 4, 1998. The consortium was announced to the public on May 20 and 21, 1998. Many companies have since joined the consortium as adopters of the technology (status as of July 11, 1998):

Ericsson	Promoter	Plantronics
Intel	Promoter	Psion
IBM	Promoter	Puma Technologies
Nokia	Promoter	Quadriga
Toshiba	Promoter	Qualcomm, Inc.
3Com		Samsung Electronics Ltd.
Axis		Siemens Forsvarsystem A/S
BreezeCOM		Symbian
Casio		Symbionics Ltd.
Cambridge consultants Ltd.		T-Span System
CETECOM GmbH		Temic Semiconductor
Cirrus Logic		TDK
Compaq Computer Corp.		TTP Communications Ltd.
Convergence Corporation		Universal Empowering Technologies
Dell Computer Corp.		VLSI Technology, Inc.
InnoLabs Corporation		Xircom
Jeeves Telecom Ltd.		
Lucent Technologies UK Ltd.		
Metricom		
Motorola		
NeoParadigm Labs, Inc.		

* The name, Bluetooth, was taken from Harald Blåtand, a Danish Viking king from the early Middle Ages.




Figure 2
Frequency-hop/time-division-duplex channel.

The Bluetooth air interface

The focus of user scenarios envisioned for first-generation products is typically on traveling business people. Portable devices that contain Bluetooth radios would enable them to leave cables and connectors at home (Box C). Before the air interface for Bluetooth could be designed, however, certain requirements had to be settled:

- The system must operate worldwide.
- The connection must support voice and data—for instance, for multimedia applications.
- The radio transceiver must be small and operate at low power. That is, the radio must fit into small, portable devices, such as mobile phones, headsets and personal digital assistants (PDA).

License-free band

To operate worldwide, the required frequency band must be available globally. Further, it must be license-free and open to any radio system. The only frequency band that satisfies these requirements is at 2.45 GHz—the Industrial-Scientific-Medical (ISM) band, which ranges from 2,400 to 2,483.5 MHz in the US and Europe (only parts of this band are available in France and Spain), and from 2,471 to 2,497 MHz in Japan. Consequently, the system can be used worldwide, given that the radio transceivers cover the frequency band between 2,400 and 2,500 MHz and that they can select the proper segment in this band.

Frequency hopping

Since the ISM band is open to anyone, radio systems operating in this band must cope with several unpredictable sources of interference, such as baby monitors, garage door openers, cordless phones and microwave ovens (the strongest source of interference). Interference can be avoided using an adaptive scheme that finds an unused part of the spectrum, or it can be suppressed by means of spectrum spreading. In the US, radios operating in the 2.45 GHz ISM band are required to apply spectrum-spreading techniques if their transmitted power levels exceed 0 dBm.

Bluetooth radios use frequency-hop (FH) spread spectrum, since this technology better supports low-cost, low-power radio implementations. Frequency-hop systems divide the frequency band into several hop channels. During a connection, radio transceivers hop from one channel to another in a pseudo-random fashion. The instantaneous (hop) bandwidth is small in frequency-hop radios, but spreading is usually obtained over the entire frequency band. This results in low-cost, narrowband transceivers with maximum immunity to interference. Occasionally, interference jams a hop channel, causing faulty reception. When this occurs, error-correction schemes in the link restore bit errors.

Channel definition

Bluetooth channels use a frequency-hop/time-division-duplex (FH/TDD) scheme (Figure 2). The channel is divided into 625 μs intervals—called slots—where a different hop frequency is used for each slot. This gives a nominal hop rate of 1,600 hops per second. One packet can be transmitted

Box C User scenarios

Three-in-one phone—use the same phone everywhere

When you are at the office, your phone functions as an intercom (no telephony charge). At home, it functions as a cordless phone (fixed-line charge). When you are on the move, it functions as a mobile phone (cellular charge).

Internet bridge—surf the Internet regardless of the connection

Use your portable PC to surf the Internet anywhere, regardless of whether you are connected wirelessly through a mobile phone (cellular) or through a wired connection (PSTN, ISDN, LAN, xDSL).

Interactive conference—connect every participant for instant data exchange

In meetings and at conferences, you can share information instantly with other participants. You can also operate a projector remotely without wire connectors.

The ultimate headset—a cordless headset keeps your hands free

Connect a headset to your mobile PC or to any wired connection and free your hands for more important tasks at the office or in your car.

Portable PC speakerphone

Connect cordless headsets to your portable PC and use it as a speakerphone regardless of whether you are in the office, your car, or at home.

Briefcase trick (hidden computing 1)

Access e-mail while your portable PC is still in

the briefcase. When your portable PC receives an e-mail message, you will be notified by your mobile phone. You can also use the phone to browse incoming e-mail and read messages.

Forbidden message (hidden computing 2)

Compose e-mail on your PC while you are on an airplane. When you land and are allowed to switch on your mobile phone, the messages are sent immediately.

Automatic synchronization (hidden computing 3)

Automatically synchronize your desktop computer, portable PC, notebook (PDA or HPC) and mobile phone. As soon as you enter the office, the address list and calendar in your notebook automatically updates the files on your desktop computer or vice versa.

Instant postcard—send photos and video clips instantly from any location

Connect a camera cordlessly to your mobile phone or to any wire-bound connection. Add comments from your mobile phone, a notebook, or portable PC and send them instantly to a recipient anywhere in the world. Suitable for professional and personal use.

Cordless desktop—connect all peripheral tools to your PC or the LAN

Connect your desktop/laptop computer cordlessly to printers, scanners and the LAN. Increase your sense of freedom through cordless mouse and keyboard connections to your PC.

per interval/slot. Subsequent slots are alternately used for transmitting and receiving, which results in a TDD scheme.

Two or more units sharing the same channel form a piconet, where one unit acts as a master, controlling traffic on the piconet, and the other units act as slaves. The frequency-hop channel is determined by the frequency-hop sequence (the order in which hops are visited) and by the phase in this sequence. In Bluetooth, the sequence is determined by the identity of the piconet master and phase is determined by the master unit's system clock (Figure 3). In order to create the master clock in the slave unit, the slave may add an offset to its own native clock. The repetition interval of the frequency-hop sequence, which is very long (more than 23 hours), is determined by the clock. If every participant on a given channel uses the same identity and clock as input to the hop-selection box, then each unit will consistently select the same hop carrier and remain synchronized. Every piconet has a unique set of master parameters which create a unique channel.

The channel makes use of several, equally spaced, 1 MHz hops. With Gaussian-shaped frequency shift keying (FSK) modulation, a symbol rate of 1 Mbit/s can be achieved. In countries where the open band is 80 MHz or broader, 79 hop carriers have been defined. In countries where the band is narrower (Japan, France, and Spain), only 23 hop carriers have been defined (Table 1). On average, the frequency-hop sequence visits each carrier with equal probability.

Packet definition

In each slot, a packet can be exchanged between the master unit and one of the slaves. Packets have a fixed format (Figure 4). Each packet begins with a 72-bit access code that is derived from the master identity and is unique for the channel. Every packet exchanged on the channel is preceded by this access code. Recipients on the piconet compare incoming signals with the access code. If the two do not match, the received packet is not considered valid on the channel and the rest of its contents are ignored. Besides packet identification, the access code is also used for synchronization and compensating for offset. The access code is very robust and resistant to interference. Correlation of the access code by recipients provides similar processing gains as direct-sequence spreading.

A header trails the access code. It contains important control information, such as a




Figure 3
Hop selection scheme: In the selection box, the master identity selects the sequence, and the clock selects the phase. Combined, they give the hop carrier to be used.

Parameters	Values
Modulation	G-FSK, $h \leq 0.35$
Peak data rate	1 Mbit/s
RF bandwidth	220 kHz ($\sim 3\text{dB}$), 1 MHz ($\sim 20\text{ dB}$)
RF band	2.4 GHz ISM band
RF carriers	23/79
Carrier spacing	1 MHz
Peak TX power	$\leq 20\text{ dBm}$

Table 1
Radio parameters.

three-bit media-access-control (MAC) address, packet type, flow control bits, bits for the automatic-retransmission-query (ARQ) scheme and a header-error-check (HEC) field (Figure 5). The header, whose length is fixed at 54 bits, is protected by a one-third rate forward-error-correction (FEC) code.

Payload may or may not trail the header. The length of the payload may vary from 0 to 2,745 bits.

To support high data rates, multi-slot packets have been defined. A packet can cover one slot, three slots, or five slots. Packets are always sent on a single-hop carrier. For multi-slot packets, the hop carrier is used as applied in the first slot. After the multi-slot packet, the channel continues on the hop as dictated by the master clock. For example, let us consider four slots: k , $k+1$, $k+2$ and $k+3$. Ordinarily, these would be associated with hop frequencies f_k , f_{k+1} , f_{k+2} and f_{k+3} . However, a three-slot packet that starts in slot k uses f_k for the entire packet. The next packet begins in slot $k+3$ and uses f_{k+3} .

Physical link definition

Two types of link have been defined for supporting multimedia applications that mix voice and data:

- synchronous connection-oriented (SCO) link;





Figure 4
Fixed packet format.

Figure 5
Header fields.



Type	Symmetric (kbit/s)	Asymmetric (kbit/s)	
DM1	108.8	108.8	108.8
DH1	172.8	172.8	172.8
DM3	256.0	384.0	54.4
DH3	384.0	576.0	86.4
DM5	286.7	477.8	36.3
DH5	432.6	721.0	57.6

Table 2
Achievable data rates (in kbit/s) on the ACL link.

- asynchronous connectionless (ACL) link. SCO links support symmetrical, circuit-switched, point-to-point connections typically used for voice. These links are defined on the channel by reserving two consecutive slots (forward and return slots) at fixed intervals.

ACL links support symmetrical or asymmetrical, packet-switched, point-to-multipoint connections typically used for bursty data transmission. Master units use a polling scheme to control ACL connections.

A set of packets has been defined for each physical link.

- For SCO links, three kinds of single-slot voice packet have been defined, each of which carries voice at a rate of 64 kbit/s. Voice is sent unprotected, but if the SCO interval is decreased, a forward-error-correction rate of 2/3 or 1/3 can be selected.
- For ACL links, 1-slot, 3-slot, and 5-slot data packets have been defined. Data can be sent either unprotected or protected by a 2/3 forward-error-correction rate. The maximum data rate—721 kbit/s in one direction and 57.6 kbit/s in the reverse direction—is obtained from an unprotected, 5-slot packet. Table 2 summarizes the data rates that can be obtained from ACL links. DMx represents x-slot, FEC-encoded data packets; DHx represents unprotected data packets.

Figure 6 depicts mixed SCO and ACL links on a piconet with one master and two slaves. Slave 1 supports an ACL link and an SCO link with a six-slot SCO interval. Slave 2 only supports an ACL link. Note: slots may be empty when no data is available.

Interference Immunity

As mentioned above, the Bluetooth radio must operate in an open band that is sub-

ject to considerable uncontrolled interference. Thus, the air interface has been optimized to deal with interference.

- Frequency hopping techniques are applied with a high hopping rate and short packet lengths (1,600 hops/s for single-slot packets). If a packet is lost, only a small portion of the message is lost.
- Packets can be protected by forward error control.
- Data packets are protected by an ARQ scheme in which lost data packets are automatically retransmitted. The recipient checks each received packet for errors. If errors are detected, it indicates this in the header of the return packet. This results in a fast ARQ scheme—delays are only one slot in duration, and only packets that have been lost need to be retransmitted.
- Voice is never retransmitted. Instead, a robust voice-encoding scheme is used. This scheme, which is based on continuous variable slope delta (CVSD) modulation, follows the audio waveform (Figure 7) and is very resistant to bit errors—errors are perceived as background noise, which intensifies as bit errors increase.

Networking


Piconets

Bluetooth units that are within range of each other can set up *ad hoc* connections. In principle, each unit is a peer with the same hardware capabilities (unlike cellular systems, there is no distinction between terminals and base stations). Two or more Bluetooth units that share a channel form a piconet. To regulate traffic on the channel, one of the participating units becomes a master of the piconet. Any unit can become a master, but by definition, the unit that establishes the piconet assumes this role. All other participants are slaves. Participants may change roles if a slave unit wants to take over the master role. Nonetheless, only one master may exist in a piconet at any time.

Every unit in the piconet uses the master identity and clock to track the hopping channel. Each unit also has its own (native), free-running clock. When a connection is established, a clock offset is added to synchronize the slave clock with the master clock. The native clock is never adjusted, however, and offsets are solely valid for the duration of the connection.

Master units control all traffic on a channel. They allocate capacity for SCO links by reserving slots. For ACL links, they use a

Figure 6
SCO and ACL links in a piconet with one master and two slaves.



polling scheme. A slave is only permitted to send in the slave-to-master slot when it has been addressed by its MAC address in the preceding master-to-slave slot. A master-to-slave packet implicitly polls the slave; that is, an ordinary traffic packet addressed to a slave polls the slave automatically. If no information to the slave is available, the master can use a POLL packet to poll the slave explicitly. POLL packets consist of an access code and header only. This central polling scheme eliminates collisions between slave transmissions.

Establishing connection

When units are not participating in a piconet, they enter standby mode, from which state they periodically listen for page messages. From the total set of 79 (23) hop carriers, a subset of 32 (16) wake-up carriers has been defined. The subset, which is chosen pseudo-randomly, is determined by the unit identity. Over the wake-up carriers, a wake-up sequence visits each hop carrier once: the sequence length is 32 (16) hops. Every 2,048 slots (1.28 s), standby units move their wake-up hop carrier forward one hop in the wake-up sequence. The native clock of the unit determines the phase of the wake-up sequence. During the listening interval, which lasts for 18 slots or 11.25 ms, the unit listens on a single wake-up hop carrier and correlates incoming signals with the access code derived from its own identity. If the correlator triggers—that is, if most of the received bits match the access code—the unit activates itself and invokes a connection-setup procedure. Otherwise, the unit returns to sleep until the next wake-up event.

Units connecting to a unit in standby mode must know the standby unit's identity and preferably its native clock

- to generate the required access code (which constitutes the paging message);
- to derive the wake-up sequence;
- to predict the phase of this sequence.

Since paging units cannot accurately know the native clock of a recipient, they must resolve the time-frequency uncertainty. They do so by transmitting the access code continuously—not only in the hop in which they expect the recipient to wake up, but also in hops before and after. For a period of 10 ms, paging units transmit the access code sequentially on several hop frequencies around the expected hop carrier. Note: the access code is only 72 bits long (72 ms). Therefore, many codes can be sent in the




Figure 7
Continuous variable slope delta (CVSD) waveform coding.

space of 10 ms. The 10 ms train of access codes on different hop carriers is transmitted repeatedly until the recipient responds or a time-out is exceeded.


When a paging unit and recipient select the same wake-up carrier, the recipient receives the access code and returns an acknowledgement. The paging unit then sends a packet containing its identity and its current clock. After the recipient acknowledges this packet, each unit uses the paging unit's parameters for hop selection—thereby establishing a piconet in which the paging unit acts as the master.

To establish a connection, the paging unit must obtain the identity of units within transmission range. Therefore, it executes an inquiry procedure: the paging unit transmits an inquiry access code (which is common to all Bluetooth devices) on the inquiry wake-up carriers. When a recipient receives the inquiry, it returns a packet containing its identity and clock—the very opposite of the paging procedure. After having gathered each response, the paging unit can then select a specific unit to page (Figure 8).

Scatternet

Users of a channel must share capacity. Although channels are 1 MHz wide, as more and more users are added, throughput per user quickly drops to less than a few tens of kbit/s. Furthermore, although the medium available bandwidth is 80 MHz in the US and Europe (slightly less in Japan, France and Spain), it cannot be used efficiently when every unit must share the same 1 MHz hop channel. Therefore, another solution was adopted.

Figure 8
Connection-establishment procedure and maximum time associated with establishing a connection.






Figure 9
A scatternet of four piconets applied to the scenario described in Figure 1.

Units that share the same area and that are within range of one another can potentially establish *ad hoc* connections between themselves. However, only those units that truly want to exchange information share the same channel (piconet). This solution permits several piconets to be created with overlapping areas of coverage. Each piconet adheres to its own hopping sequence through the 80 MHz medium. The channel in each piconet hops pseudo-randomly over the carriers in the 80 MHz band. The users in each piconet have only a 1 MHz hop channel at their disposal.

A group of piconets is called a scatternet. Aggregate and individual throughput of users in a scatternet is much greater than when each user participates on the same piconet with a 1 Mbit/s channel. Additional gains are obtained by statistically multiplexing hop channels and by reusing channels. The 1 MHz hop channel in any given

piconet need only be shared by users of that piconet. Because individual piconets hop differently, different piconets can simultaneously use different hop channels. Consequently, units in one piconet do not share their 1 MHz channel with units in another piconet. The aggregate throughput (the total throughput accumulated over all piconets) increases as more piconets are added. Collisions do occur, however, when two piconets use the same hop channel simultaneously. As the number of piconets increases, performance in the frequency-hop system degrades gracefully. Simulations of a scatternet consisting of 10 piconets indicate that reduction in throughput per piconet is less than 10%. In the scatternet, the radio medium is shared; in a piconet, the channel and information are shared.

Since every piconet uses the same bandwidth, each shares the 80 MHz in an average sense. Provided they select different hop channels, however, no two piconets must simultaneously share the same 1 MHz channel.

Let us assume there are 100 users. If each belonged to the same network, all 100 users would have to share the same 1 MHz channel. Thus, average throughput per user would be 10 kbit/s and aggregate throughput would be 1 Mbit/s. However, if not everyone wanted to talk to each other, we could split the piconet into independent piconets. For example, if the users separated themselves into groups of five, then we could create 20 independent piconets. With only five users sharing the 1 MHz hop channel, throughput per user increases to 200 kbit/s and aggregate throughput increases to 20 Mbit/s. Obviously, this assumes ideal conditions, where no two piconets select the same hop channel at the same time. In reality—because the piconets hop independently—collisions will occur, reducing effective throughput. Nonetheless, the final throughput obtained from multiple piconets exceeds that of a single piconet.

The maximum number of units that can actively participate on a single piconet is eight: one master and seven slaves. The MAC address in the packet header, which is used to distinguish each unit, is limited to three bits. Figure 9 illustrates the scatternet approach applied to the scenario shown in Figure 1.




Figure 10
Hop selection in inter-piconet communication.

Inter-piconet communication

Different piconets adhere to different frequency-hop sequences and are controlled

by different masters. If a hop channel is temporarily shared by independent piconets, packets can be distinguished by the access codes that precede them—access codes are unique for each piconet. Piconets are uncoordinated and hop independently: synchronization of different piconets is not permitted in the unlicensed ISM band. Nonetheless, units may participate in different piconets on a time-division-multiplexing (TDM) basis. That is, a unit can sequentially participate in different piconets, provided it is active in only one piconet at a time.

Inter-piconet communication is achieved by selecting the proper master identity and clock offset in order to synchronize with the channel of the desired piconet (Figure 10). A corresponding set of identity and clock offsets is available for each piconet. Whenever a unit enters a piconet, it adjusts the clock offset to account for minor drifts between the master clock and the unit's native clock. A unit can thus act as a slave in several piconets. When leaving one piconet for another, a slave informs the current master that it will be unavailable for a predetermined period. During its absence, traffic on the piconet between the master and other slaves continues as usual.

A master unit can also periodically jump to another piconet and act as a slave (were it to act as a master in the new piconet, that piconet would have the same channel parameters as the "old" piconet—therefore, by definition the two would be indiscernible). When a master unit leaves a piconet, all traffic on the piconet is suspended until it returns.

Figure 11 shows a slave participating in two piconets. Piconet X consists of master X and slaves A_x and B_x . Piconet Y consists of master Y and slaves A_y , B_y , and D_y . Slave C_{xy} participates in piconets X and Y. The clock of each unit is also shown. Positive offset (indicated by blue) or negative offset (indicated by red) has been added for synchronization with the master clock. Slave C_{xy} contains a native clock and two offsets for the master units X and Y respectively.

Authentication and encryption

To ensure user protection and information secrecy, the system must provide security measures that are appropriate for a peer environment—that is, each unit in Bluetooth must implement authentication and en-




Figure 11
Participation of slave C_{xy} in two piconets.

ryption algorithms in the same way.

A base level encryption has been specified that is well suited to silicon implementation, and an authentication algorithm has been specified that provides a level of security for devices lacking in processing capabilities. Future support for ciphering algorithms will be backward-compatible.

The main security features are:

- a challenge-response routine—for authentication;
- stream cipher—for encryption;
- session key generation—session keys can be changed at any time during a connection.

Three entities are used in the security algorithms: the Bluetooth unit address, which is a public entity; a private user key, which is a secret entity; and a random number, which is different for each new transaction. As described above, the Bluetooth address can be obtained through an inquiry procedure. The private key is derived during initialization and is never disclosed. The random number is derived from a pseudo-random process in the Bluetooth unit.

Conclusion

Bluetooth is a system for providing local wireless connectivity between portable devices. It is particularly suitable for *ad hoc* networking. The air interface has been optimized to provide maximum immunity against interference in the 2.45 GHz band. The system is supported by several leading manufacturers of personal computers and telecommunications equipment. The first consumer products to support Bluetooth are expected to appear on the market around year-end 1999.